

The Complexity of Sensing by Point Sampling

Yan-Bin Jia, *Carnegie Mellon University, Pittsburgh, PA*

Michael Erdmann, *Carnegie Mellon University, Pittsburgh, PA*

In assembly tasks it is often necessary to recognize parts arriving via a conveyor belt or a parts feeder at some robot work cell. Generally the parts feeder will have reduced the number of possible poses of the parts to a small finite set. In order to distinguish between the remaining poses of the parts some simple sensing or probing operation may be used. In this paper we consider the problem of finding the minimum number of sensing points required to distinguish between a finite set of polygonal shapes. For instance, we might imagine embedding a series of point light detectors in a feeder tray. Then we would be interested in the question “What is the minimum number of light detectors that can fully distinguish between all the possible shapes?” Or we might imagine a set of mechanical probes that touches the feeder at a finite number of predetermined points. Then we would ask “What are the minimum number of probing points and where should the probes be located in order to distinguish all the possible shapes?” We address these questions in this paper.

*Intuitively, each sensing point can be regarded as a binary bit that has two values ‘contained’ and ‘not contained’. So the robot senses a shape by reading out the binary representation of the shape, that is, by checking which points are contained in the shape and which are not. The formalized sensing problem: Given n polygons with a total of m edges in the plane, locate the fewest points such that each polygon contains a distinct subset of points in its interior. We show that this problem is equivalent to an NP-complete set-theoretic problem introduced as *Discriminating Set*. By a reduction to *Hitting Set* (and hence to *Set Covering*), an $O(n^2m^2)$ approximation algorithm is presented to solve the sensing problem with a ratio of $2 \ln n$. Based on a reverse reduction, we prove that one can use an algorithm for *Discriminating Set* with ratio $c \log n$ to construct an algorithm for *Set Covering* with ratio*

*$c \log n + O(\log \log n)$. Thus approximating *Discriminating Set* exhibits the same hardness as that of approximating *Set Covering* recently shown in [24] and [4]; this result implies that the ratio $2 \ln n$ is asymptotically optimal unless $\text{NP} \subset \text{DTIME}(n^{\text{poly} \log n})$. Finally we analyze the complexity of subproblems of *Discriminating Set*, based on their relationship to a generalization of *Independent Set* called *3-Independent Set*.*

1 Introduction

One of the fundamental tasks in automatic assembly is for robots to efficiently determine the positions and orientations, termed the *poses*, of the individual parts to be assembled. The geometric shapes of these parts are designed early in the manufacturing process, so they are known in advance. Often the possible poses in which a part settles on the assembly table are of a small number, either reduced by a parts feeder or limited by the mechanical constraints imposed by a sequence of planned manipulations. (See Erdmann and Mason’s tray-tilting method [14] and Brost’s squeeze-grasp method [6] for examples of the latter case.) These facts together allow the implementation of effective sensing mechanisms, which usually take the form of simple and fast hardware systems coupled with efficient geometric algorithms [7]. The efficiency of such sensing mechanisms depends on both the time cost of the physical operations and the time complexity of the algorithms involved. Consequently, minimizing one or both of these two factors has become an important aspect of sensor design.

In order to illustrate the goals of this paper, consider a polygonal part resting on a horizontal assembly table. The table is bounded by vertical fences at its bottom-left corner, as shown in Figure 1. Pushing the part towards that corner will eventually cause the part to settle in one of the 12 stable poses listed in the figure.

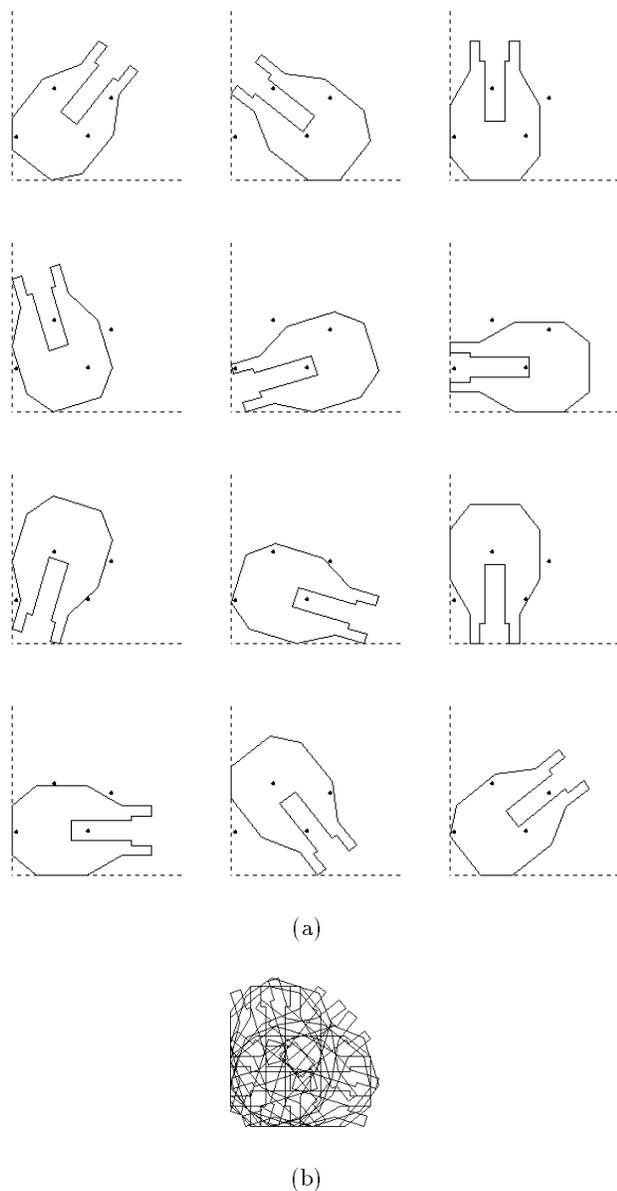


Figure 1: Sensing by point sampling. (a) The 12 possible stable poses of an assembly part after pushing, along with 4 sampling points (optimal by Lemma 1) found by our implementation of the approximation algorithm to recognize these poses, where dashed line segments are fences perpendicular to the plane. (b) The planar subdivision formed by these poses which consists of 610 regions.

(Note to reach a stable pose both fences must be in contact with some vertices of the part while at least

one fence must be in contact with no less than two vertices.) In order to distinguish between these 12 poses, the robot has marked 4 points on the table beforehand, so it can infer the pose from which marks are covered by the part and which are not¹.

The above “shape recovery” method is named *sensing by point sampling*, as a loose analogy to the reconstruction of band limited functions by sampling on a regular grid in signal processing. To save the expense of sampling, the robot wants to mark as few points as possible. The problem: *How to compute a minimum set of points to be marked so that parts of different types and poses can be distinguished from each other by this method?*

1.1 Related Work

Natarajan [25] examined a similar strategy of detecting the orientations of polygonal and polyhedral objects with an analysis of the numbers of sensors sufficient and necessary for the task. More recent related work includes [5] and [2]. [5] shows that the problem of deciding whether k line probes are sufficient to distinguish a convex polygon from a collection of n convex polygons is NP-complete. This result is very similar to our Theorem 2. A variation of the line-probing result in [5] would give us the point sampling result of Theorem 2. [2] proves a similar result as well, namely that the problem of constructing a decision tree of minimum height to distinguish among n polygons using point probes is NP-complete. This result holds even if all the polygons are convex. [2] also exhibits a greedy approximation algorithm for constructing such a decision tree. This result is similar to our approximation algorithm of Section 4, with a similar ratio bound. The difference is that our greedy algorithm seeks to minimize the total number of probe points rather than the tree height.²

¹This can be implemented in multiple ways, such as placing light detectors in the table, probing at the points, or if the robot has a vision system, taking a scene image and checking the corresponding pixel values.

²It is easy to give an example for which a minimum height decision tree uses more than minimum number of total probes, while a decision tree with minimum number of total probes does not attain the minimum height. Consider the problem of discriminating sets $\{a, b, a', e'\}$, $\{a, a'\}$, $\{b, b', d'\}$, $\{c, b'\}$, $\{d, c'\}$ and \emptyset which can be viewed as probing a collection of polygons by the later transformation

Closely related work includes the research by Romanik and others on geometric testability (see, for example, [27], [28], and [1]). Their research develops strategies for verifying a given polygon using a series of point probes. Moreover, the research examines the testability of more general geometric objects, such as polyhedra, and develops conditions that determine whether a class of objects is (approximately) testable.

A number of researchers have looked at the problem of determining or distinguishing objects using finger probes. Finger probing is closely related to sensing by point sampling, as indicated by our discussion of [5]. For a more extensive survey of probing problems and solutions see the paper by Skiena [29].

There would seem to be connections between our work and the concept of VC-dimension often used in learning theory. For instance, in this paper we develop the notion of a “discriminating set” to distinguish different polygons. The concept of a discriminating set bears some resemblance to the idea of shattered sets associated with VC-dimension. However, discriminating sets and shattered sets are different. A minimum discriminating set is the smallest set of points that uniquely identifies every object in a set of objects, whereas VC-dimension is the size of the largest set of points shattered by the set of objects. Thus, the VC-dimension of a finite class gives a lower bound on the size of a minimum discriminating set. For dense polygon distributions, the two cardinalities will be the same, namely $\log n$, where n is the number of polygons. For sparsely distributed polygons, the two cardinalities are different. For instance, the VC-dimension can be 1, while the minimum discriminating set has size $n - 1$. See Figure 2.

Finally, the work described in this paper is part of our larger research goal to understand the information requirements of robot tasks. Related work includes the sensor design methodology of Erdmann [13] and the information invariants of Donald *et al.* [11]. [13] proposes a method for designing sensors based on the particular manipulation task at hand. The resulting sensors satisfy a minimality property with respect to the given task goal and the available robot actions. [11] inves-

technique in the proof of Theorem 2. The decision tree using minimum probes a, b, c, d always has height 4, and the decision tree using probes a', b', c', d', e' can achieve minimum height 3.

tigates the relationship between sensing, action, distributed resources, communication paths, and computation, in the solution of robot tasks. That work provides a method for comparing disparate sensing strategies, and thus for developing minimal or redundant strategies, as desired.

1.2 The Formal Problem

Consider n simple polygons P_1, \dots, P_n in the plane, not necessarily disjoint from each other. We wish to locate the minimum number of points in the plane such that no two polygons P_i and P_j , $i \neq j$, contain exactly the same points. In order to avoid ambiguities in sensing, we require that none of the located points lie on any edge of P_1, \dots, P_n . The planar subdivision formed by P_1, \dots, P_n divides the plane into one unbounded region, some bounded regions outside P_1, \dots, P_n , called the “holes”, and some bounded regions inside. (For example, the 12 polygons in Figure 1(a) form the subdivision in Figure 1(b) which consists of 610 regions, none of which is a hole.) Immediately we make two observations: (1) Points on the edges of the subdivision or in the interior of the unbounded region or in a “hole” do not need to be considered as locations; (2) for each bounded (open) region inside some polygon only one point needs to be considered.

Let Ω denote the set of bounded regions in the subdivision which are contained in at least one of P_1, \dots, P_n . Each polygon P_i , $1 \leq i \leq n$, is partitioned into one or more such regions; we write $\omega \sqsubseteq P_i$ when a region ω is contained in polygon P_i . A *region basis* for polygons P_1, \dots, P_n is a subset $\Delta \subseteq \Omega$ such that

$$\{\omega \mid \omega \in \Delta \text{ and } \omega \sqsubseteq P_i\} \neq \{\omega \mid \omega \in \Delta \text{ and } \omega \sqsubseteq P_j\},$$

for $1 \leq i \neq j \leq n$; that is, each P_i contains a distinct collection of regions from Δ . A region basis Δ^* of minimum cardinality is called a *minimum region basis*. Thus the problem of sensing by point sampling becomes the problem of finding a minimum region basis Δ^* . We will call this problem *Region Basis* and focus on it throughout the paper. The following lemma gives the upper and lower bounds for the size of such Δ^* .

Lemma 1 *A minimum region basis Δ^* for n polygons P_1, \dots, P_n satisfies $\lceil \log n \rceil \leq |\Delta^*| \leq n - 1$.*

Proof. To verify the lower bound $\lceil \log n \rceil$, note that each of the n polygons must contain a distinct subset

of Δ^* ; so $n \leq 2^{|\Delta^*|}$, the cardinality of the power set 2^{Δ^*} .

To verify the upper bound $n - 1$, we incrementally construct a region basis Δ of size at most $n - 1$. This construction is similar to Natarajan's Algorithm 2 [25]. Initially, $\Delta = \emptyset$. If $n > 1$, without loss of generality, assume P_1 has the *smallest area*. Then there exists some region $\omega_1 \in \Omega$ outside P_1 . Split $\{P_1, \dots, P_n\}$ into two *nonempty* subsets, one including those P_i containing ω_1 and the other including those not; and add ω_1 into Δ . Recursively split the resulting subsets in the same way, and at each split, add into Δ its defining region (as we did with ω_1) if this region is not already in Δ , until every subset eventually becomes a singleton. The Δ thus formed is a region basis. Since there are $n - 1$ splits in total and each split adds at most one region into Δ , we have $|\Delta| \leq n - 1$. \square

Figure 2 gives two examples for which $|\Delta^*| = \lceil \log n \rceil$ and $|\Delta^*| = n - 1$ respectively. Therefore these two bounds are tight.

We can view all the bounded non-hole regions as elements of Ω , and all the polygons P_1, \dots, P_n as subsets of Ω . Then a region basis Δ is a subset of Ω that can discriminate subsets P_1, \dots, P_n by intersection. Hence the Region Basis problem can be rephrased as: Find a subset of Ω of minimum size whose intersections with any two subsets P_i and P_j , $1 \leq i \neq j \leq n$, are not equal. The general version of this set-theoretic problem, in which Ω stands for an arbitrary finite set and P_1, \dots, P_n stand for arbitrary subsets of Ω , we call *Discriminating Set*. We have thus reduced Region Basis to Discriminating Set, and the former problem will be solved once we solve the latter one.

Let us analyze the amount of computation required for the geometric preprocessing to reduce Region Basis to Discriminating Set. Let m be the total *size* of P_1, \dots, P_n , i.e., the sum of the number of vertices each polygon has; trivially $m \geq 3$. Then the planar subdivision these polygons define has at most s vertices, where $3 \leq s \leq \binom{m}{2}$. By Euler's relation on planar graphs, the number of regions and the number of edges are upper bounded by $2s - 4$ and $3s - 6$ respectively. So we can construct the planar subdivision either in time $O(m \log m + s)$ using an optimal algorithm for intersecting line segments by Chazelle and Edelsbrunner [8], or in time $O(s \log m)$ using a simpler plane sweep version

by Nievergelt and Preparata [26]. To obtain the set of regions each polygon contains, we only need to traverse the portion of the subdivision bounded by that polygon, which takes time $O(s)$. It follows that the reduction to Discriminating Set can be done in time $O(m \log m + ns)$, or $O(nm^2)$ in the worst case.

Here is a short summary of the structure of the paper: Section 2 proves the NP-completeness of Discriminating Set; based on this result, Section 3 establishes an equivalence between Discriminating Set and Region Basis, hence proving the latter problem NP-complete; Section 4 presents an $O(n^2 m^2)$ approximation algorithm for Region Basis with ratio $2 \ln n$ and shows that further improvements on this ratio are hard; and Section 5 closes up with a complexity analysis of various subproblems of Discriminating Set, along with the definition of a family of related NP-complete problems called k -Independent Sets. We have implemented our approximation algorithm and have tested it on both real data taken from mechanical parts and random data extracted from the arrangements of random lines. The algorithm works very well in practice.

2 Discriminating Set

Given a collection C of subsets of a finite set X , suppose we want to identify these subsets just from their *intersections* with some subset $D \subseteq X$. Thus D must have distinct intersection with every member of C , that is,

$$D \cap S \neq D \cap T, \quad \text{for all } S, T \in C \text{ and } S \neq T.$$

We call such a subset D a *discriminating set* for C with respect to X . From a different point of view, each element $x \in D$ can be regarded as a binary "bit" that, to represent any subset $S \subseteq X$, gives value '1' if $x \in S$ and value '0' otherwise. In such a way D represents an encoding scheme for subsets in C .

Below we show that the problem of finding a minimum discriminating set is NP-complete. As usual, we consider the decision version of this minimization problem:

DISCRIMINATING SET (D-SET)

Let C be a collection of subsets of a finite set X and $l \leq |X|$ a non-negative integer. Is there a discriminating set $D \subseteq X$ for C such that $|D| \leq l$?

The Complexity of Sensing by Point Sampling

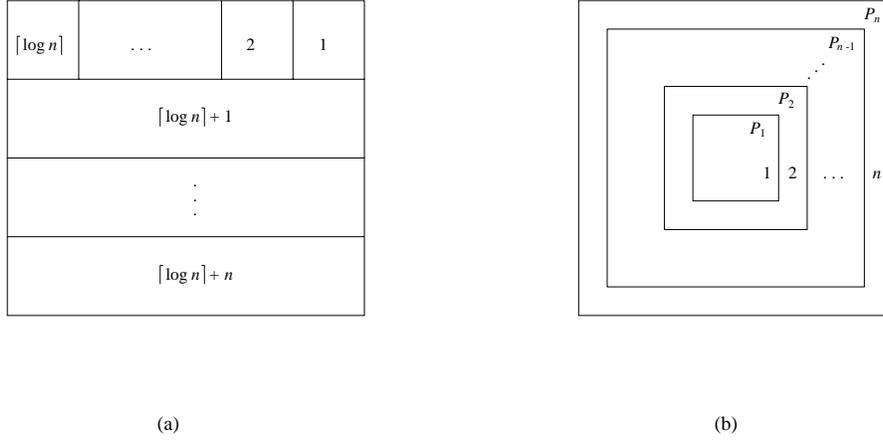


Figure 2: Two examples whose minimum region basis sizes achieve the lower bound $\lceil \log n \rceil$ and the upper bound $n - 1$ respectively. Bounded regions in the examples are labelled with numbers. (a) For $1 \leq i \leq n$ polygon P_i is defined to be the boundary of the union of regions $\lceil \log n \rceil + 1, \dots, \lceil \log n \rceil + i$, and all regions k with $1 \leq k \leq \lceil \log n \rceil$ such that the k th bit of the binary representation (radix 2) for $i - 1$ is 1. Thus $\Delta^* = \{1, 2, \dots, \lceil \log n \rceil\}$. (b) The polygons P_1, \dots, P_n contain each other in increasing order: $\Delta^* = \{2, 3, \dots, n\}$.

Our proof of the NP-completeness for D-Set uses a reduction from Vertex Cover (VC) which determines if a graph $G = (V, E)$ has a *cover* of size not exceeding some integer $l \geq 0$, i.e., a subset $V' \subseteq V$ that, for each edge $(u, v) \in E$, contains either u or v . The reduction is based on a key observation, that for any three finite sets S_1, S_2 and S_3 ,

$$S_1 \cap S_2 \neq S_1 \cap S_3 \quad \text{if and only if} \quad S_1 \cap (S_2 \Delta S_3) \neq \emptyset,$$

where ‘ Δ ’ denotes the operation of symmetric difference, i.e., $S_2 \Delta S_3 = (S_2 \setminus S_3) \cup (S_3 \setminus S_2)$.

Theorem 1 *Discriminating Set is NP-complete.*

Proof. That D-Set \in NP is trivial.

Next we establish $\text{VC} \leq_{\text{P}} \text{D-Set}$, that is, there exists a polynomial-time reduction from VC to D-Set. Let $G = (V, E)$ and integer $0 < l \leq |V|$ be an instance of VC. We need to construct a D-Set instance (X, C) such that the collection C of subsets of X has a discriminating set of size l' or less if and only if G has a vertex cover of size l or less.

The construction uses the component design technique described by Garey and Johnson [16]. It’s rather natural for us to begin by including every vertex of G in set X , and assigning each edge $e = (u, v)$ a subset $S(e)$ in C which contains at least u and v ; in other

words, we have $V \subset X$ and

$$\{u, v\} \subset S(e) \in C, \quad \text{for all } e = (u, v) \in E.$$

In order to ensure that any discriminating set D for C contains at least one of u and v from subset $S(e)$, we add an *auxiliary* subset A_e into C that consists of some new elements *not* in V , and in the meantime define

$$S(e) = \{u, v\} \cup A_e.$$

Hence $S(e) \Delta A_e = \{u, v\}$; and $D \cap \{u, v\} \neq \emptyset$ follows directly from $D \cap S(e) \neq D \cap A_e$. Since any discriminating set D' for $\{A_e \mid e \in E\}$ can also distinguish between $S(e_1)$ and $S(e_2)$, and between $S(e_1)$ and A_{e_2} , for any $e_1, e_2 \in E$ and $e_1 \neq e_2$, D' unioned with a vertex cover for G becomes a discriminating set for C . Conversely, every discriminating set D for C can be split into a discriminating set for $\{A_e \mid e \in E\}$ and a vertex cover for G .

The $m = |E|$ auxiliary subsets should be constructed in a way such that we can easily determine the size of their minimum discriminating sets in order to set up the entire D-Set instance. There is a simple way: We introduce m elements $a_1, a_2, \dots, a_m \notin V$ into X , and define subsets A_e , for $e \in E$, to be

$$\{a_1\}, \{a_2\}, \dots, \{a_m\},$$

where the order of mapping does not matter. It’s clear that there are m minimum discriminating sets for the above subsets: $\{a_1, \dots, a_m\} \setminus \{a_i\}$, $1 \leq i \leq m$.

Setting $l' = l + m - 1$, we have completed our construction of the D-Set instance as

$$\begin{aligned} X &= V \cup \{a_1, \dots, a_m\}; & a_1, \dots, a_m &\notin V; \\ C &= \{S(e) \mid e \in E\} \cup \{A_e \mid e \in E\}. \end{aligned}$$

The construction can be carried out in time $O(|V| + |E|)$. We omit the remaining task of verifying that G has a vertex cover of size at most l if and only if C has a discriminating set of size at most $l + m - 1$. \square

One thing about this proof is worthy of note. All subsets in C constructed above have at most three elements. This reveals that D-Set is still NP-complete even if $|S| \leq 3$ for all $S \in C$, a stronger assertion than Theorem 1. The subproblem where all $S \in C$ have $|S| \leq 1$ is obviously in P, for an algorithm can simply count $|C|$ in linear time and then answer “yes” if $l \geq |C| - 1$ and “no” if $0 \leq l < |C| - 1$. For the remaining case in which all $S \in C$ have $|S| \leq 2$, we will prove in Section 5 that the NP-completeness still holds. However, the proof will be a bit more involved than the one we just gave under no restriction on $|S|$.

At the end of this section, we give a problem that is equivalent to D-Set:

ROW-DIFFERING SUBMATRIX

Given an $m \times n$ matrix A of 0's and 1's and integer $0 \leq l \leq n$, is there an $m \times l$ matrix B formed by l columns of matrix A such that no two rows of B are identical?

3 Region Basis

Now that we have shown the NP-completeness of D-Set, the minimum region basis cannot be computed in polynomial time through the use of an efficient algorithm for D-Set, because no such algorithm would exist unless $P = NP$. This conclusion, nevertheless, leads us to conjecture that the minimization problem Region Basis is also NP-complete. Again we consider the decision version:

REGION BASIS (RB)

Given n polygons P_1, \dots, P_n and integer $0 \leq l \leq n - 1$, does there exist a region basis Δ for the planar subdivision Ω formed by P_1, \dots, P_n such that $|\Delta| \leq l$?

The condition $0 \leq l \leq n - 1$ above is necessary because we already know from Lemma 1 that a minimum region basis has size at most $n - 1$.

Consider a mapping \mathcal{F} from the set of RB instances to the set of D-Set instances that maps regions to elements and polygons to subsets in a one-to-one manner. Every RB instance is thus mapped into an equivalent D-Set instance, as pointed out in Section 1. We claim that \mathcal{F} is not onto. Suppose \mathcal{F} were onto. Then the elements of each subset in a D-Set instance must correspond to regions in some RB instance. The union of these regions must be a polygon, and this polygon must map to the subset given in the D-Set instance. However, this is not always possible. Consider a D-Set instance generated from a nonplanar graph such that each edge is a subset containing its two vertices as only elements. No RB instance can be mapped to such a D-Set instance. For if there were such an RB instance, the geometric dual of the planar subdivision it defines would contain a planar embedding for the original nonplanar graph. This is an impossibility, hence we have a contradiction.

Thus the set of RB instances constitutes a *proper* subset of the set of D-Set instances; in other words, RB is isomorphic to a subproblem of D-Set. Therefore, the NP-completeness of RB does not follow directly from that of D-Set established earlier. Fortunately, however, D-Set has an equivalent subproblem which is isomorphic to a subproblem of RB under \mathcal{F} . That isomorphism provides us with the NP-completeness of Region Basis.

Theorem 2 Region Basis is NP-complete.

Proof. That $RB \in NP$ is easy to verify, based on the fact mentioned in Section 1 that the number of regions in the planar subdivision is at most quadratic in the total size of the polygons.

Let (X, C) be a D-Set instance, where

$$\begin{aligned} X &= \{x_1, x_2, \dots, x_m\}; \\ C &= \{S_1, S_2, \dots, S_n\} \subseteq 2^X. \end{aligned}$$

Without loss of generality, we make two assumptions

$$\bigcup_{i=1}^n S_i = X \quad \text{and} \quad \bigcap_{i=1}^n S_i = \emptyset,$$

because elements contained in none of the subsets or contained in all subsets can always be removed

from any discriminating set of (X, C) . Now add in a new element $a \notin X$ and consider the D-Set instance $(X \cup \{a\}, C')$, where $C' = \{S_i \cup \{a\} \mid 1 \leq i \leq n\}$. Clearly $(X \cup \{a\}, C')$ and (X, C) have the same set of irreducible discriminating sets³ and hence they are considered equivalent.

The planar subdivision defined by the constructed RB instance for $(X \cup \{a\}, C')$ takes the configuration shown in Figure 3(a): A rectangular region is divided by a horizontal line segment into two identical regions of which the bottom region is named $\omega(a)$; the top region is further divided, this time by vertical line segments, into $2m - 1$ identical regions of which the odd numbered ones, from left to right, are named $\omega(x_1), \dots, \omega(x_m)$ respectively. Remove those $m - 1$ unnamed regions on the top. For $1 \leq i \leq n$ define polygon P_i to be the boundary of the union of all regions $\omega(x)$, $x \in S_i \cup \{a\}$. It should be clear that P_i is indeed a polygon; and the two assumptions guarantee that P_1, \dots, P_n form the desired subdivision. Note the subdivision consists of $m + 1$ rectangular regions and $4m + 2$ vertices. All can be computed in time $\Theta(m)$, given the coordinates of the four vertices of the bounding rectangle. Thus the reduction takes time $\Theta(\sum_{i=1}^n |S_i|)$.

It is clear that C has a discriminating set of size l or less if and only if there is a region basis of the same size for P_1, \dots, P_n . Hence we have proved the NP-completeness of RB. \square

The above proof implies that we may regard Discriminating Set and Region Basis as equivalent problems. Note that the polygons P_1, \dots, P_n in Figure 3 are not convex; will Region Basis become P when all the polygons are convex? This question is answered by the following corollary.

Corollary 1 *Region Basis remains NP-complete even if all the polygons are convex.*

Proof. Same as the proof of Theorem 2 except that we use the planar subdivision shown in Figure 3(b). (The vertices of the subdivision partition an imaginary circle (dotted) into $2n$ equal arcs.) \square

³A discriminating set D is said to be *irreducible* if no subset $D' \subset D$ can be a discriminating set.

4 Approximation

Sometimes we can derive a polynomial-time approximation algorithm for the NP-complete problem at hand from some existing approximation algorithm for another NP-complete problem, by reducing one problem to the other. In fewer cases, where the reduction *preserves* the solutions, namely, every instance of the original problem and its reduced instance have the same set of solutions, any approximation algorithm for the reduced problem together with the reduction will solve the original problem. The problem to which we will reduce Discriminating Set is Hitting Set:

HITTING SET

Given a collection C of subsets of a finite set X , find a minimum *hitting set* for C , i.e., a subset $H \subseteq X$ of minimum cardinality such that $H \cap S \neq \emptyset$ for all $S \in C$?

Karp [20] shows Hitting Set to be NP-complete by a reduction from Vertex Cover. The reducibility from D-Set to Hitting Set follows a key fact we observed when proving Theorem 1: The intersections of a finite set D with two finite sets S and T are not equal if and only if D intersects their symmetric difference $S \Delta T$. Given a D-Set instance, the corresponding Hitting Set instance is constructed simply by replacing all the subsets with their *pairwise* symmetric differences. Thus every discriminating set of the original D-Set instance is also a hitting set of the constructed instance, and vice versa.

The approximability of Hitting Set can be studied through another problem, Set Covering:

SET COVERING

Given a collection C of subsets of a finite set X , find a minimum *cover* for X , i.e., a subcollection $C' \subseteq C$ of minimum size such that $\bigcup_{S \in C'} S = X$?

This problem is also shown to be NP-complete by Karp by a reduction from Exact Cover by 3-Sets [20]. A greedy approximation algorithm for this problem due to Johnson [18] and Lovász [23] guarantees to find a cover \hat{C} for X with ratio

$$\frac{|\hat{C}|}{|C^*|} \leq H(\max_{S \in C} |S|) \quad \text{or simply} \quad \frac{|\hat{C}|}{|C^*|} \leq \ln |X| + 1,$$

where C^* is a minimum cover and $H(k) = H_k = \sum_{i=1}^k \frac{1}{i}$, known as the k th harmonic number. The algorithm works by selecting, at each stage, a subset from

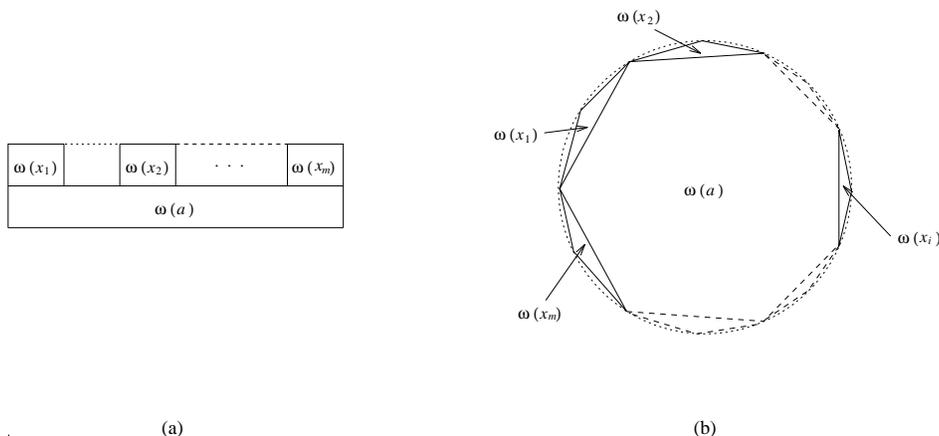


Figure 3: Two reductions from Discriminating Set to Region Basis.

C that covers the most remaining uncovered elements of X . We refer the reader to [9] for a general analysis of the greedy heuristic for Set Covering.

Hitting Set and Set Covering are *duals* to each other—the roles of set and element in one problem just get switched in the other. More specifically, let a Hitting Set instance consist of some finite set X and a collection C of its subsets; its dual Set Covering instance then consists of a set \bar{C} and a collection of its subsets \bar{X} where

$$\bar{C} = \{\bar{S} \mid S \in C\} \quad \text{and} \quad \bar{X} = \{\bar{x} \mid x \in X\},$$

and where each subset \bar{x} is defined as^{4,5}

$$\bar{x} = \{\bar{S} \mid S \in C \text{ and } S \ni x\}.$$

Intuitively speaking, the element $x \in X$ “hits” the subset $S \in C$ in the original instance if and only if the subset \bar{x} “covers” the element $\bar{S} \in \bar{C}$ in the dual instance. Thus it follows that $H \subseteq X$ is a hitting set for C if and only if $\bar{H} = \{\bar{x} \mid x \in H\}$ is a cover for \bar{C} . Hence the corresponding greedy algorithm for Hitting Set selects at each stage an element

⁴According to this definition, $\bar{x} = \bar{y}$ may hold for two different elements $x \neq y$. In this case only one subset is included in \bar{X} .

⁵This definition also establishes the duality between D-Set and a known NP-complete problem called Minimum Test Set (see [16]). Given a collection of subsets of a finite set, Minimum Test Set asks for a minimum subcollection such that exactly one from each pair of distinct elements is contained in some subset from this subcollection.

that “hits” the most remaining subsets. It is clear that the approximation ratio for Hitting Set becomes $H(\max_{x \in X} |\{S \mid S \in C \text{ and } S \ni x\}|)$ or $\ln|C| + 1$.⁶

As a short summary, the greedy heuristic on a Discriminating Set instance (X, C) works by finding a hitting set for the instance $(X, \{S \Delta T \mid S, T \in C\})$. Since an element can appear in at most $\lfloor \frac{n^2}{4} \rfloor$ such pairwise symmetric differences, where $n = |C|$, the approximation ratio attained by this heuristic is $\ln \lfloor \frac{n^2}{4} \rfloor + 1 < 2 \ln n$. The same ratio is attained for Region Basis by the heuristic that selects at each step a region discriminating the most remaining pairs of polygons, where n is now the number of polygons.

The greedy algorithm for Set Covering (dually for Hitting Set) can be carefully implemented to run in time $O(\sum_{S \in C} |S|)$ [10]. The reduction from a D-Set instance (X, C) to a Hitting Set instance takes time $O(|C|^2 \max_{S \in C} |S|)$. Combining the time complexity of the geometric preprocessing in Section 1, we can easily verify that Region Basis can be solved in time $O(nm^2 + n^2m^2) = O(n^2m^2)$, where n and m are the number and size of polygons respectively.

In the remainder of this section we establish the hardness of approximating D-Set and hence Region Basis. Both problems allow the same approximation ratio since the reductions from one to another do not

⁶Kolaitis and Thakur [22] syntactically define a class of NP-complete problems with logarithmic approximation algorithms that contains Set Covering and Hitting Set, and show that Set Covering is complete for the class.

The Complexity of Sensing by Point Sampling

change the number of subsets (or polygons) in an instance. First we should note that the ratio bound $H(\max_{S \in C} |S|)$ of the greedy algorithm for Set Covering is actually tight; an example that makes the algorithm achieve this ratio for arbitrarily large $\max_{S \in C} |S|$ is given in [18].

Next we present a reverse reduction from Hitting Set to D-Set to show that an algorithm for D-Set with approximation ratio $c \log n$ can be used to obtain an algorithm for Hitting Set with ratio $c \log n + O(\log \log n)$, where $c > 0$ is any constant and n is the number of subsets in an instance. Afterwards, we will apply some recent results on the hardness of approximating Set Covering (and thus Hitting Set).

Lemma 2 *For any $c > 0$, if $c \log n$ is the approximation ratio of Discriminating Set, then Hitting Set can be approximated with ratio $c \log n + O(\log \log n)$.*

Proof. Suppose there exists an algorithm \mathcal{A} for D-Set with approximation ratio $c \log n$. Let (X, C) be an arbitrary instance of Hitting Set, where $C = \{S_1, \dots, S_n\} \subseteq 2^X$, and let $n = |C|$. To construct a D-Set instance, we first make $f(n)$ isomorphic copies $(X_1, C_1), \dots, (X_{f(n)}, C_{f(n)})$ of (X, C) such that $X_i \cap X_j = \emptyset$ for $1 \leq i \neq j \leq f(n)$. Here f is an as yet undetermined function of n upper bounded by some polynomial in n . Now consider the enlarged Hitting Set instance $(X', C') = (\bigcup_{i=1}^{f(n)} X_i, \bigcup_{i=1}^{f(n)} C_i)$. Every hitting set H' of (X', C') has $H' = \bigcup_{i=1}^{f(n)} H_i$, where H_i is a hitting set of (X_i, C_i) , $1 \leq i \leq n$; so from H' we can obtain a hitting set H of (X, C) with $|H| \leq |H'|/f(n)$ merely by taking the smallest one of $H_1, \dots, H_{f(n)}$.

Next we introduce a set A consisting of new elements $a_1, a_2, \dots, a_{\log(nf(n))} \notin X'$; and for $1 \leq i \leq nf(n)$ define auxiliary sets A_i :

$$A_i = \{a_j \mid 1 \leq j \leq \log(nf(n)) \text{ and the } j\text{th bit of the binary representation of } i-1 \text{ is } 1\}.$$

It is not hard to see $\{a_1, \dots, a_{\log(nf(n))}\}$ must be a subset of any discriminating set for $A_1, \dots, A_{nf(n)}$; therefore it is the minimum one for these auxiliary sets. The constructed D-Set instance is then defined to be (X'', C'') , where

$$\begin{aligned} X'' &= X_1 \cup \dots \cup X_{f(n)} \cup \{a_1, a_2, \dots, a_{\log(nf(n))}\}; \\ C'' &= \{T \cup A_{(i-1)n+j} \mid T \in C_i \text{ and } T \cong S_j\} \cup \{A_1, \dots, A_{nf(n)}\}. \end{aligned}$$

It is easy to verify that every discriminating set of (X'', C'') is the union of A and a hitting set of (X', C') .

Now run algorithm \mathcal{A} on the instance (X'', C'') and let D be the discriminating set found. Then

$$\frac{|D|}{|D^*|} \leq c \log(|C''|) = c \log(2nf(n)),$$

where D^* is a minimum discriminating set. From the construction of (X'', C'') we know that $D = H_1 \cup \dots \cup H_{f(n)} \cup A$ and $D^* = H_1^* \cup \dots \cup H_{f(n)}^* \cup A$, where for $1 \leq i \leq n$, H_i and H_i^* are some hitting set and some minimum hitting set of (X_i, C_i) , respectively. Let H_k satisfy $|H_k| = \min_{i=1}^{f(n)} |H_i|$ and thus let $H \cong H_k$ be a hitting set of (X, C) . Also, let H^* with $|H^*| = |H_1^*| = \dots = |H_{f(n)}^*|$ be a minimum hitting set of (X, C) . Then

$$\begin{aligned} \frac{|D|}{|D^*|} &= \frac{\sum_{i=1}^{f(n)} |H_i| + |A|}{\sum_{i=1}^{f(n)} |H_i^*| + |A|} \\ &\geq \frac{f(n) \cdot |H| + \log(nf(n))}{f(n) \cdot |H^*| + \log(nf(n))}. \end{aligned}$$

Combining the two inequalities above generates:

$$\begin{aligned} \frac{|H|}{|H^*|} &\leq c \log(2nf(n)) + \frac{(c \log(2nf(n)) - 1) \cdot \log(nf(n))}{f(n) \cdot |H^*|} \\ &\leq c \log(2nf(n)) + \frac{(c \log(2nf(n)) - 1) \cdot \log(nf(n))}{f(n)} \\ &= c \log n + \left[c + c \log f(n) + \frac{(c \cdot (1 + \log n + \log f(n)) - 1) \cdot (\log n + \log f(n))}{f(n)} \right]. \end{aligned}$$

Setting $f(n) = \log^2 n$, all terms in the brackets can be absorbed into $O(\log \log n)$ after simple manipulations on asymptotics [17]; thus we have

$$\frac{|H|}{|H^*|} \leq c \log n + O(\log \log n).$$

□

Though Set Covering has been extensively studied since the mid 70's, essentially nothing on the hardness of approximation was known until very recently. The results of [3] imply that no polynomial approximation scheme exists unless $P = NP$. Based on recent results from interactive proof systems and probabilistically checkable proofs and their connection to approximation, several asymptotic improvements on the hardness of approximating Set Covering have been

made. In particular, Lund and Yannakakis [24] showed that Set Covering cannot be approximated with ratio $c \log n$ for any $c < \frac{1}{4}$ unless $\text{NP} \subset \text{DTIME}(n^{\text{poly} \log n})$; Bellare et al. [4] showed that approximating Set Covering within any constant is NP-complete, and approximating it within $c \log n$ for any $c < \frac{1}{8}$ implies $\text{NP} \subset \text{DTIME}(n^{\log \log n})$. Based on their results and by Lemma 2, we conclude on the same hardness of approximating D-Set and Region Basis:

Theorem 3 *Discriminating Set and Region Basis cannot be approximated by a polynomial-time algorithm with ratio bound $c \log n$ for any $c < \frac{1}{4}$ unless $\text{NP} \subset \text{DTIME}(n^{\text{poly} \log n})$, or for any $c < \frac{1}{8}$ unless $\text{NP} \subset \text{DTIME}(n^{\log \log n})$.*

Following the above theorem, the ratio $2 \ln n \approx 1.39 \log n$ of the greedy algorithm for D-Set remains asymptotically optimal if NP is not contained in $\text{DTIME}(n^{\text{poly} \log n})$.

5 More on Discriminating Set

Now let's come back to where we left the discussion on the subproblems of D-Set in Section 2; it has not been settled whether D-Set remains NP-complete when every subset S in the collection C satisfies $|S| \leq 2$. We now prove that this subproblem is NP-complete.

Here we look at a special case of this subproblem, namely, a “subsubproblem” of D-Set, subject it to two restrictions: (1) $\emptyset \in C$ and (2) $|S| = 2$ for all nonempty subsets $S \in C$. Let's call this special case *0-2 D-Set*. If 0-2 D-Set is proven to be NP-complete, so will be the original subproblem.

It's quite intuitive to understand a 0-2 D-Set instance in terms of a graph $G = (V, E)$, where $V = X$, the finite set of which every $S \in C$ is a subset, and

$$E = \left\{ (u, v) \mid \{u, v\} \in C \right\}.$$

In other words, each element of the set X corresponds to a vertex in G while each subset, except \emptyset , corresponds to an edge. Clearly this correspondence from all 0-2 D-Set instances to all graphs is one-to-one. Since any discriminating set D for C has

$$D \cap S \neq D \cap \emptyset = \emptyset, \quad \text{for all } S \in C \text{ and } S \neq \emptyset,$$

D must be a vertex cover for G . Let $d(u, v)$ be the *distance*, i.e., the length of the shortest path, between

vertices u, v in G (or ∞ if u and v are disconnected). A *3-independent set* in G is a subset $I \subseteq V$ such that $d(u, v) \geq 3$ for every pair $u, v \in I$. The following lemma captures the dual relationship between a discriminating set for C and a 3-independent set in G .

Lemma 3 *Let X be a finite set and C a collection of \emptyset and two-element subsets of X . Let $G = (X, E)$ be a graph with $E = \{(u, v) \mid \{u, v\} \in C\}$. Then a subset $D \subseteq X$ is a discriminating set for C if and only if $X \setminus D$ is a 3-independent set in G .*

Proof. Let D be a discriminating set for C . Assume there exist two distinct elements (vertices) $u, v \in X \setminus D$ such that $d(u, v) < 3$. We immediately have $(u, v) \notin E$ since D must be a vertex cover in G ; so $d(u, v) = 2$. Hence there is a third vertex, say w , that is connected to both u and v ; furthermore, $w \in D$ holds since the edges (u, w) and (v, w) must be covered by D . But now we have $D \cap \{u, w\} = D \cap \{v, w\} = \{w\}$, a contradiction to the fact that D is a discriminating set.

Conversely, suppose $X \setminus D$ is a 3-independent set in G , for some subset $D \subseteq X$. Then D must be a vertex cover. Suppose it is not a discriminating set for C . Then there exist two distinct subsets $S_1, S_2 \in C$ such that $D \cap S_1 = D \cap S_2 = \{w\}$, for some $w \in S$. Writing $S_1 = \{u, w\}$ and $S_2 = \{v, w\}$, we have $d(u, v) = 2$; but in the meantime $u, v \in X \setminus D$. A contradiction again. \square

This lemma tells us that the NP-completeness of 0-2 D-Set, and therefore of our remaining open subproblem of D-Set, follows if we can show the NP-completeness of *3-Independent Set*. 3-Independent Set is among a family of problems defined, for all integers $k > 0$, as follows:

k -INDEPENDENT SET (k -IS)

Given a graph $G = (V, E)$ and an integer $0 < l \leq |V|$, is there a *k-independent set* of size at least l , that is, is there a subset $I \subseteq V$ with $|I| \geq l$ such that $d(u, v) \geq k$ for every pair $u, v \in I$?

Thus 2-IS is the familiar NP-complete Independent Set problem. We will see in Appendix A that every problem in this family for which $k > 3$ is also NP-complete. To avoid too much divergence from 0-2 D-Set, let's focus on 3-IS only here.

Lemma 4 *3-Independent Set is NP-complete.*

The Complexity of Sensing by Point Sampling

Proof. It is trivial that 3-IS \in NP. To show NP-hardness, we reduce Independent Set (2-IS) to 3-IS. Let $G = (V, E)$ and $0 < l \leq |V|$ form an instance of Independent Set. A graph G' is constructed from G in two steps. In the first step, we introduce a “midvertex” $\alpha_{u,v}$ for each edge $(u, v) \in E$, and replace this edge with two edges $(u, \alpha_{u,v})$ and $(\alpha_{u,v}, v)$. In the second step, an edge is added between every two midvertices that are adjacent to the same original vertex. More formally, we have defined $G' = (V', E')$ where

$$\begin{aligned} V' &= V \cup \left\{ \alpha_{u,v} \mid (u, v) \in E \right\}; \\ E' &= \left\{ (\alpha_{u,v}, u) \mid (u, v) \in E \right\} \cup \\ &\quad \left\{ (\alpha_{u,v}, \alpha_{u,w}) \mid (u, v) \neq (u, w) \in E \right\}. \end{aligned}$$

Two observations are made about this construction. First, it has the property that $d'(u, v) = d(u, v) + 1$ holds for any pair of vertices $u, v \in V$, where d and d' are the two distance functions in G and G' respectively. This equality can be verified by contradiction. Next, if $(u, v) \in E$, then any two midvertices $\alpha_{u,x}$ and $\alpha_{v,y}$ have

$$\begin{aligned} d'(\alpha_{u,x}, \alpha_{v,y}) &\leq d'(\alpha_{u,x}, \alpha_{u,v}) + d'(\alpha_{u,v}, \alpha_{v,y}) \leq 2; \\ d'(\alpha_{u,x}, v) &= d'(\alpha_{u,x}, \alpha_{u,v}) + d'(\alpha_{u,v}, v) \leq 2; \\ d'(\alpha_{v,y}, u) &= d'(\alpha_{v,y}, \alpha_{u,v}) + d'(\alpha_{u,v}, u) \leq 2. \end{aligned}$$

Note strict ‘<’s appear in above three inequalities when $x = v$ or $y = u$, and in the first inequality when $x = y$. It is not difficult to see that the whole reduction can be done in time $O(|V|^3)$. Figure 4 illustrates an example of the reduction.

We claim that G has an independent set I of size at least l if and only if G' has a 3-independent set I' of the same size. Suppose I with $|I| \geq l$ is an independent set in G . Then I is also a 3-independent set in G' . This follows from our first observation. Conversely, suppose I' with $|I'| \geq l$ is a 3-independent set in G' . Then the set I , produced by replacing each midvertex $\alpha_{u,v} \in I'$ with either u or v , is an independent set in G . To see this, assume there exists two vertices $u, v \in I$ such that $d(u, v) = 1$. Thus $d'(u, v) = d(u, v) + 1 = 2 < 3$; so either u or v , or both, must have replaced some midvertices in I' . Let $s, t \in I'$ be the two vertices corresponding to u and v before the replacement respectively; that is, $s = u$ or $\alpha_{u,x}$ and $t = v$ or $\alpha_{v,y}$ for some $x, y \in V$. According to our second observation, we always have $d'(s, t) \leq 2$. Thus we have reached a

contradiction, since $s, t \in I'$. That $|I'| = |I| \geq l$ is easy to verify in a similar way. \square

Combining Lemmas 3 and 4, we have the NP-completeness of 0-2 D-Set; this immediately resolves the complexity of our remaining subproblem of D-Set:

Theorem 4 *D-Set remains NP-complete even if $|S| \leq 2$ for all $S \in C$.*

6 Experiments

For geometric preprocessing, we implemented the plane sweep algorithm by Nievergelt and Preparata [26]. We modified the original algorithm so that the containing polygons of each swept region are maintained and propagated along during the sweeping.⁷ The greedy approximation algorithm for Set Covering was implemented with a linked list to attain the running time $O(\sum_{S \in C} |S|)$. All code was written in Common Lisp and was run on a Sparcstation IPX.

We discuss simulation results on random polygons. These simulations empirically study how the number of sampling points varies with the “density” of polygons in the plane. The results suggest that the point sampling approach is most effective at sensing polygonal objects that have highly overlapping poses. Experiments on a Zebra robot are underway and the results will be presented in the near future.

6.1 Simulation Results

To generate random polygons, we precomputed an arrangement of a large number (such as 100) of random lines using Edelsbrunner and Guibas’s topological sweeping algorithm [12]. A random polygon was extracted as the first “valid” cycle occurring in a random walk on this line arrangement, after being scaled to some random perimeter. By “valid” we mean that the number of vertices in the cycle was no less than some small random integer. This constraint was introduced merely to allow a proper distribution of polygons of

⁷This implementation has the same worst-case running time as a different version described in Section 1 which obtains the containment information by traversing the planar subdivision after the sweeping. But the implementation version is usually more efficient in practice.

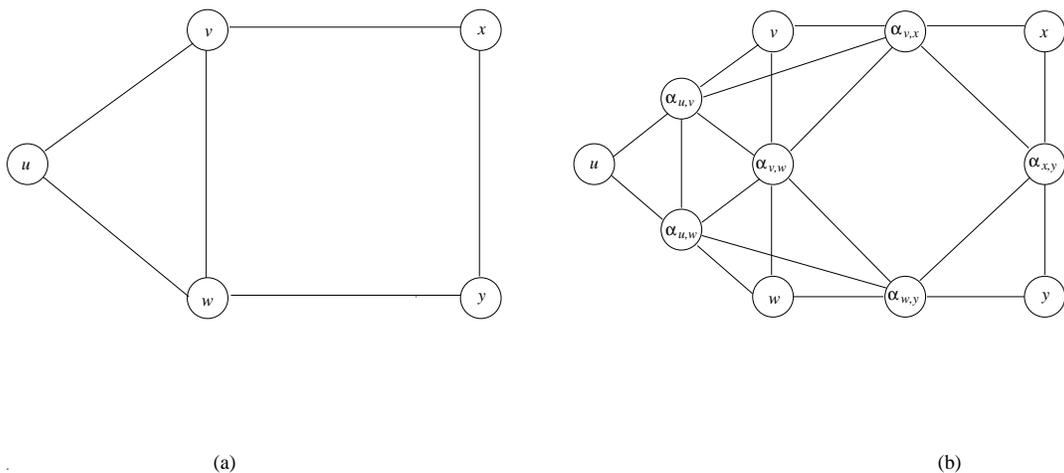


Figure 4: An example of reduction from Independent Set to 3-IS. (a) An instance of Independent Set. (b) The constructed 3-IS instance: α vertices are added to increase the distances between the original vertices by exactly one.

various sizes, for otherwise triangles and quadrilaterals would be generated with high probabilities according to our observations. In a sample run, a group of 1000 polygons generated (by this method) from an arrangement of 100 random lines had sizes in the range 3–30, with mean 5.545 and standard deviation 3.225.

All random polygons (or all random poses of a single polygon) in a test were bounded by some square, so that the “density”, i.e., the degree of overlap, of these polygons mainly depended on their number as well as on the ratio between their average area and the size of the bounding square. Since polygons were generated randomly, the average area could be viewed as approximately proportional to the square of the average perimeter. The configuration of each polygon, say P , was assumed to obey a “uniform” distribution inside the square. More specifically, the orientation of P was first randomly chosen from $[0, 2\pi)$; the position of P was then randomly chosen from a rectangle inside the square consisting of all feasible positions at that orientation.⁸

To be robust against sensor noise, the sampling point of every region in the region basis was selected as the center of a maximum inscribed circle in that region. In other words, this sampling point had the maximum

⁸If the diameter of P is greater than the width of the square, then not every orientation is necessarily feasible. However, this situation was avoided in our simulations.

distance to the polygon bounding that region. It is not difficult to see that such a point must occur at a vertex of the generalized Voronoi diagram inside the polygon, also called its *internal skeleton* or *medial axis function*.⁹ Also for sensing robustness, regions with area less than some threshold were not considered at the stage of region basis computation.¹⁰ Though this thresholding traded off the completeness of sampling, it almost never resulted in the failure of finding a region basis once the threshold was properly set.

The first two groups of six tests gave a sense of the number of sampling points required when polygons are sparsely distributed in the plane. The results are summarized in Table 1. Every test in group (a) was conducted on distinct, i.e., non-congruent, random polygons with perimeters between $\frac{1}{4}$ and $\frac{3}{4}$ of the width of the bounding square; every test in group (b) was conducted on distinct poses of a single polygon with

⁹The construction of the internal skeleton of a polygon is a special case of the construction of the generalized Voronoi diagram for a set of line segments, for which $O(n \log n)$ algorithms are given in [15], [21], and [30]. Since the maximum region size for a region basis turned out to be very small in the simulations, we only implemented an $O(n^4)$ brute force algorithm.

¹⁰We thresholded on the region area rather than the radius of a maximum inscribed circle merely to avoid the inefficient computation on the latter for all the regions in the planar subdivision.

The Complexity of Sensing by Point Sampling

# polys	# regions	# sampling points
50	320	31
60	500	37
70	594	41
80	783	46
90	973	47
100	1422	51

(a)

# polys	# regions	# sampling points
50	362	36
60	609	34
70	741	34
80	1061	39
90	1125	49
100	1643	61

(b)

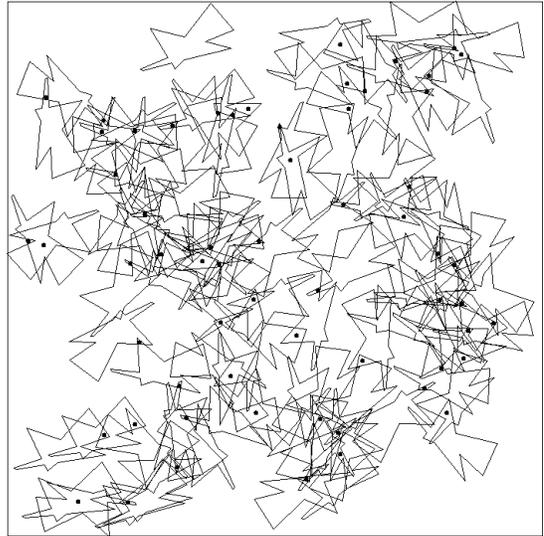
Table 1: Tests on sampling sparsely distributed random polygons/poses. The twelve tests were divided into two groups: (a) All polygons in each test were distinct, with perimeters between $\frac{1}{4}$ and $\frac{3}{4}$ times the width of the bounding square. (b) All polygons in each test represented distinct random poses of a same polygon. The polygon perimeter was uniformly $\frac{5}{8}$ times the side length of the square for all six tests in the group.

perimeter equal to $\frac{5}{8}$ of the width of the square. The scene of the last test from each group is displayed in Figure 5.

Without any surprise, the number of sampling points found were *around half* of the number of polygons, for all twelve tests in Table 1. This supports the fact that, for n sparsely distributed polygons in the plane, the minimum number of sampling points turns out to be $\Theta(n)$. As we can see from Figure 5, in such a situation every polygon intersects at most a few, or more precisely, a constant number of, other polygons. In other words, the number of polygon pairs distinguishable by any single region in the planar subdivision is $\Theta(n)$; but there are $\lfloor \frac{n^2}{4} \rfloor$ such pairs in total! Thus sensing by point sampling is inefficient in a situation with a large number of sparsely distributed polygons.



(a)



(b)

Figure 5: Sampling 100 sparsely distributed random polygons/poses. (a) The scene of the last test from group (a) in Table 1: There are 1422 regions in the planar subdivision and 51 sampling points (drawn as dots) to discriminate the 100 polygons. (b) The scene of the last test from group (b) in Table 1: There are 1643 regions in the planar subdivision and 61 sampling points to discriminate the 100 poses.

The next two groups of six tests were on polygons much more densely distributed in the plane, and the results are given in Table 2. In these two groups of

# polys	# regions	# sampling points
25	776	8
30	998	8
35	1270	12
40	1759	12
45	2153	11
50	2678	13

(c)

# polys	# regions	# sampling points
10	264	4
20	1121	7
25	1781	6
30	2796	7
35	3655	9
40	4995	9

(d)

Table 2: Tests on sampling densely distributed random polygons/poses, divided into two groups (c) and (d). The width of the bounding square was reduced to $\frac{1}{4}$ times the width of the square used in groups (a) and (b) in Table 1. In group (c) all polygons in each test were distinct with perimeters in the range between $\frac{1}{2}$ and 2 times the width of the bounding square. In group (d) all polygons in each test were distinct poses of the same polygon as in Figure 5(b).

tests, we used a bounding square with side length only $\frac{1}{4}$ of the width of the one used in test groups (a) and (b). Every test in group (c) was conducted on distinct polygons with perimeters in the range $\frac{1}{2}$ –2 times the side length of the bounding square. All tests in group (d) were distinct poses of the same polygon used in the last test of group (b). Again the scene of the last test from each group is shown in Figure 6.

All twelve tests in groups (c) and (d) except the last one in group (c) found sampling points *at most twice* the lower bound $\lceil \log n \rceil$, while the first test in group (d) found exactly $\lceil \log n \rceil$ sampling points. The data in group (d) were more densely distributed than

the data in group (c) in that every two poses intersected each other. Since an extremely dense distribution of polygons may cause numerical instabilities in the plane sweep algorithm, smaller numbers of polygons were tested in these two groups than were tested in groups (a) and (b). The results of these two groups of tests show that the sampling strategy is very applicable to sensing densely distributed polygons.

Acknowledgments

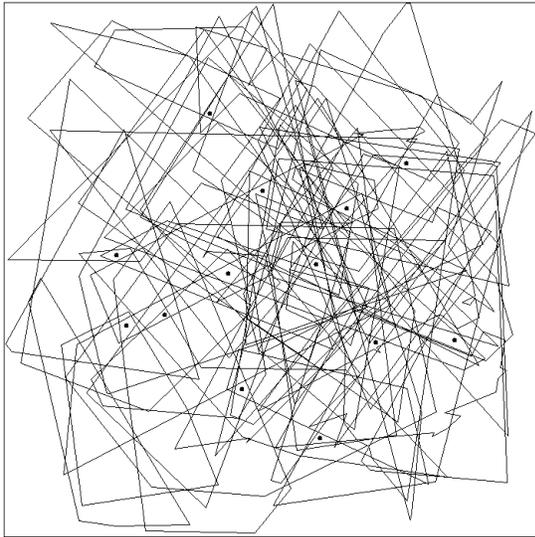
Support for this research was provided in part by Carnegie Mellon University, and in part by the National Science Foundation through the following grants: NSF Research Initiation Award IRI-9010686, NSF Presidential Young Investigator award IRI-9157643, and NSF Grant IRI-9213993.

Many thanks to David S. Johnson for pointing to the results of [24] on the hardness of approximating Set Covering and for looking into the status of Discriminating Set and k -Independent Sets, which we found neither in the catalog [16] nor in the NP-Completeness Columns of *Journal of Algorithms* starting from [19]. Also thanks to Somesh Jha for his valuable suggestions on the proof of Lemma 4, and to Bruce Donald for his valuable comments and reading of this paper.

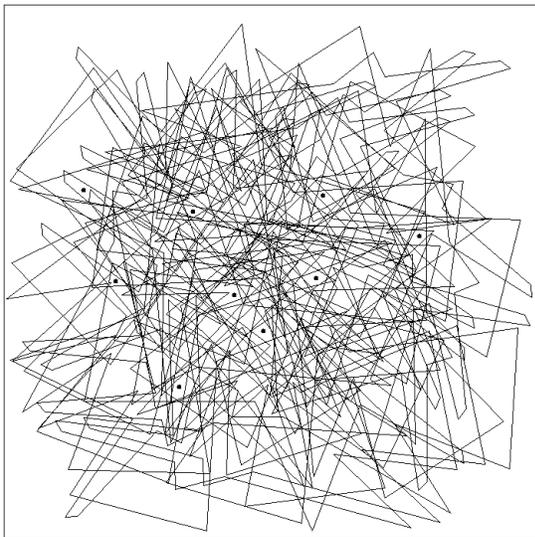
After presentation of this paper at the workshop we learned that similar results to our Theorem 2 and our approximation algorithm had appeared in [5], [27], [1], and [2]. We are very grateful to Jean-Daniel Boissonnat, Mark Overmars, Anil Rao, and Kathleen Romanik for pointing out these papers. We are particularly grateful to Kathleen Romanik for numerous interesting discussions on shape discrimination and shattered sets, and for her reading of this paper.

References

- [1] Esther M. Arkin, Patrice Belleville, Joseph S. B. Mitchell, David Mount, Kathleen Romanik, Steven Salzberg, and Diane Souvaine. Testing Simple Polygons. In *Proceedings of the Fifth Canadian Conference on Computational Geometry*, August 5–9, 1993.
- [2] Esther Arkin, Hank Meijer, Joseph S. B. Mitchell, David Rappaport, and Steven S. Skiena. Decision Trees for Geometric Models. In *Proceedings of the*



(c)



(d)

Figure 6: Sampling densely distributed polygons. The bounding square has width $\frac{1}{4}$ times the width of the one shown in Figure 5. (c) The scene of the last test from group (c) in Table 2: There are 50 distinct polygons which form a planar subdivision with 2678 regions, and which can be discriminated by 13 sampling points. (d) The scene of the last test from group (d) in Table 2: There are 40 distinct poses of the polygon from Figure 5(b), which form a planar subdivision with 4955 regions, and which can be discriminated by 9 sampling points.

Ninth Annual Symposium on Computational Geometry, pages 369–378, May 19–21, 1993.

- [3] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and intractability of approximation problems. In *Proceedings of the 33rd IEEE Symposium on Foundations of Computer Science*, pages 14–23, 1992.
- [4] M. Bellare, S. Goldwasser, G. Lund, and A. Russell. Efficient probabilistically checkable proofs and applications to approximation. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, pages 294–304, 1993.
- [5] Patrice Belleville and Tom C. Shermer. Probing Polygons Minimally is Hard. *Computational Geometry: Theory and Applications*, 2(5):255–265, 1993.
- [6] R.C. Brost. Automatic grasp planning in the presence of uncertainty. *International Journal of Robotics Research*, 7(1):3–17, 1988.
- [7] John F. Canny and Kenneth Y. Goldberg. A “RISC” paradigm for industrial robotics. Technical Report ESRC 93-4, University of California at Berkeley, 1993.
- [8] Bernard Chazelle and Herbert Edelsbrunner. An optimal algorithm for intersecting line segments in the plane. *Journal of the ACM*, 39(1):1–54, 1992.
- [9] V. Chvátal. A greedy heuristic for the set-covering problem. *Mathematics of Operations Research*, 4(3):233–235, 1979.
- [10] Thomas H. Cormen, Charles E. Leiserson, and Ronald L. Rivest. *Introduction to Algorithms*. McGraw-Hill, 1990.
- [11] Bruce R. Donald, Jim Jennings, and Daniela Rus. Towards a theory of information invariants for cooperating autonomous mobile robots. *Sixth International Symposium on Robotics Research*, October 1–5, 1993, Hidden Valley, Pennsylvania.
- [12] Herbert Edelsbrunner and Leonidas J. Guibas. Topologically sweeping an arrangement. In *Proceedings of the 18th Annual ACM Symposium on Theory of Computing*, pages 389–403. ACM Press, 1986.
- [13] Michael Erdmann. Understanding action and sensing by designing action-based sensors. *Sixth*

- International Symposium on Robotics Research*, October 1–5, 1993, Hidden Valley, Pennsylvania.
- [14] Michael Erdmann and Matt Mason. An exploration of sensorless manipulation. *IEEE Journal of Robotics and Automation*, 4(4):369–379, 1988.
- [15] Steven Fortune. A sweepline algorithm for Voronoi diagrams. *Algorithmica*, 2:153–174, 1987.
- [16] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Free and Company, 1979.
- [17] Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete Mathematics: A Foundation for Computer Science*. Addison-Wesley, 1989.
- [18] David S. Johnson. Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences*, 9:256–278, 1974.
- [19] David S. Johnson. The NP-completeness column: an ongoing guide. *Journal of Algorithms*, 4:393–405, 1981.
- [20] Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller and James W. Thatcher, editors, *Complexity of Computer Computations*. Plenum Press, 1972.
- [21] David G. Kirkpatrick. Efficient computation of continuous skeletons. In *Proceedings of the 20th Annual Symposium on Foundations of Computer Science*, pages 18–27, 1979.
- [22] Phokion G. Kolaitis and Madhukar N. Thakur. Approximation properties of NP minimization classes. In *Proceedings of the 6th Conference on Structure in Complexity Theory*, pages 353–366, 1991.
- [23] Lászo Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13:383–390, 1975.
- [24] Carsten Lund and Mihalis Yannakakis. On the hardness of approximation minimization problems. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, pages 286–293. ACM Press, 1993.
- [25] B.K. Natarajan. On detecting the orientation of polygons and polyhedra. In *Proceedings of the 3rd Annual Symposium on Computational Geometry*, pages 146–152, 1987.
- [26] J. Nievergelt and F. P. Preparata. Plane-sweep algorithms for intersecting geometric figures. *Communications of the ACM*, 25(10):739–747, 1982.
- [27] Kathleen Romanik. Approximate Testing Theory. Ph.D. Thesis, Technical Report CS-TR-2947/UMIACS-TR-92-89, University of Maryland at College Park, August 1992.
- [28] Kathleen Romanik and Steven Salzberg. Testing Orthogonal Shapes. *Proceedings of the Fourth Canadian Conference on Computational Geometry*, pages 216–222, August 10–14, 1992.
- [29] Steven S. Skiena. Problems in Geometric Probing. *Algorithmica*, 4:599–605, 1989.
- [30] Chee K. Yap. An $O(n \log n)$ algorithm for the Voronoi diagram of the set of simple curve segments. *Discrete & Computational Geometry*, 2:365–393, 1987.

Appendix

A k -Independent Sets

We extend Lemma 4 to all k -IS with $k > 3$: They are NP-complete as well. The proof we will present is indeed a generalization of the proof of Lemma 4; it will again construct a k -IS instance with graph G' from an instance of Independent Set with graph G by local replacement. In the proof, each vertex v in G will be replaced by a simple path \mathcal{P}_v of fixed length (depending only on k) that has v in the middle and an equal number of auxiliary vertices on each side; and each edge (u, v) will be replaced by four edges connecting the two end vertices on \mathcal{P}_u with the two end vertices on \mathcal{P}_v , either directly or through a “midvertex”. More intuitively speaking, all shortest paths between pairs of vertices in G , if they exist, get elongated in G' to such a degree that (1) (u, v) is an edge in G if and only if the distance between vertices u and v in G' is less than k ; (2) any two vertices u' and v' in G' with a distance of at least k can be easily mapped to two *nonadjacent* vertices in G . The first condition ensures that any given independent set in G will be a k -independent set in G' , while the second condition ensures the construction of an independent set in G from any given k -independent set in G' .

Lemma 5 *k -Independent Set is NP-complete for all integers $k > 3$.*

Proof. Given an instance of Independent Set as a graph $G = (V, E)$ and a positive integer $l \leq |V|$, a k -IS instance is constructed by two consecutive substitutions. A path

$$\mathcal{P}_v = \begin{cases} v_{k-3} \dots v_1 v v_2 \dots v_{k-2}, & \text{if } k \text{ even;} \\ v_{k-4} \dots v_1 v v_2 \dots v_{k-3}, & \text{if } k \text{ odd,} \end{cases}$$

first substitutes for vertex $v \in V$, where v_1, \dots, v_{k-3} (and v_{k-2} when k is even) are auxiliary vertices. And then a set of four edges

$$E_{u,v} = \begin{cases} \left\{ (u_{k-3}, v_{k-3}), (u_{k-3}, v_{k-2}), \right. \\ \quad \left. (u_{k-2}, v_{k-3}), (u_{k-2}, v_{k-2}) \right\}, & \text{if } k \text{ even;} \\ \left\{ (u_{k-4}, \alpha_{u,v}), (u_{k-3}, \alpha_{u,v}), \right. \\ \quad \left. (v_{k-4}, \alpha_{u,v}), (v_{k-3}, \alpha_{u,v}) \right\}, & \text{if } k \text{ odd,} \end{cases}$$

substitute for each edge $(u, v) \in E$, where $\alpha_{u,v}$ is an introduced midvertex. Figure 7 shows two subgraphs after applying the above substitutions on edge $(u, v) \in E$, for k even and odd respectively.

We can easily verify that, for any pair of vertices x on \mathcal{P}_u and y on \mathcal{P}_v , both k even and odd, we have $d'(x, y) \leq d'(u, v) = k - 1 < k$ if $(u, v) \in E$, where d' is the distance function defined on G' . On the other hand, if $(u, v) \notin E$, we have $d'(u, v) \geq k$ when k is even and $d'(u, v) \geq k + 1$ when k is odd. Thus an independent set I in G is also a k -independent set in G' . Conversely, suppose I' with $|I'| \geq l$ is a k -independent set in G' . We substitute $u \in V$ for every auxiliary vertex $u_i \in I'$ on path \mathcal{P}_u , and u or v for every midvertex $\alpha_{u,v} \in I'$ when k is odd. Let I be the set after this substitution. It needs to be shown that I is an independent set in G and $|I| = |I'| \geq l$. This is obvious for the case that k is even. When k is odd, however, the situation is a bit more complicated due to the possible occurrences of those α vertices in I' . We observe, for any $\alpha_{u,v}, \alpha_{u',v'}$, and x on path \mathcal{P}_w where $u, v, u', v', w \in V$,

$$\begin{aligned} d'(\alpha_{u,v}, x) &\leq \frac{k-3}{2} + 3 < k, & \text{if } w = u \text{ or } v, \\ & & \text{or } (u, w) \in E, \\ & & \text{or } (v, w) \in E; \\ d'(\alpha_{u,v}, \alpha_{u',v'}) &\leq 4 < k, & \text{if } (u, u') \in E. \end{aligned}$$

In fact, these two conditions guarantee that I is an independent set in G and $|I| = |I'|$, which we leave for the reader to verify.

The reduction can be done in time $O(k|V| + |E|)$, which reduces to $O(|V| + |E|)$ if k is treated as a constant, in contrast to the time $O(|V|^3)$ required for the reduction from Independent Set to 3-IS. This time reduction is due to the fact that midvertices corresponding to the same vertex in V no longer have edges between each other. \square

Since 1-IS can be easily solved by comparing $|V|$ and l , we are now ready to sum up the complexity results on this family of problems in the following theorem.

Theorem 5 *k -Independent Set is in P if $k = 1$ and NP-complete for all $k \geq 2$.*

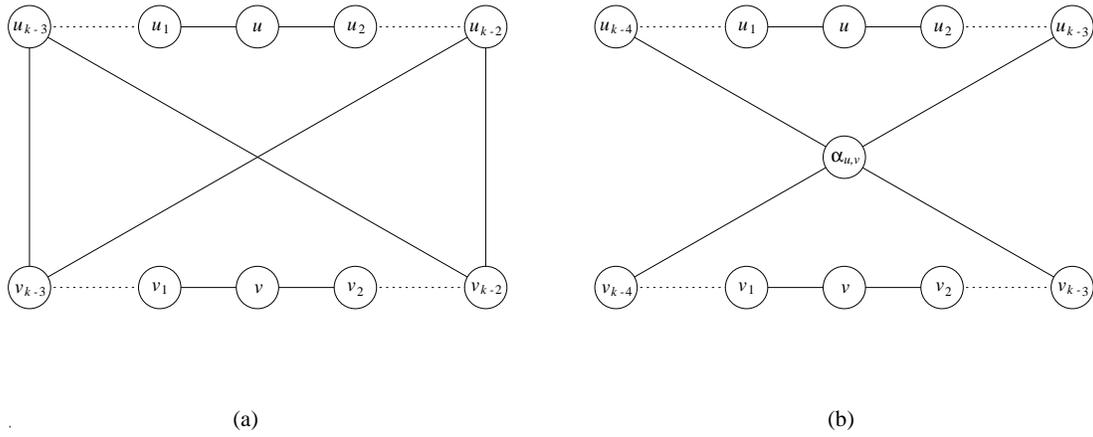


Figure 7: Two subgraphs resulting from the described substitutions performed on edge $(u, v) \in E$ for the cases that (a) k is even; and (b) k is odd.